

REMARKS

Claims 1, 2, 4-9 and 12-17 are currently active.

The Examiner has rejected Claims 1, 2, 4-6, 8 and 9 as being unpatentable over Lee in view of Dighe. Applicants respectfully traverse this rejection.

Referring to Lee, there is disclosed a method for dynamic connection management in integrated communication networks. Lee teaches in a connection-oriented communication network, the transfer information between two end-users is accomplished by network functions that select and allocate the resources along an acceptable path. The logical association between the communicating end-users is referred to as a call. The chain of associated network resources that support a call is referred to as a connection. Connection management is a network function that is responsible for setting up, maintaining, and taking down connections. Each call request is issued with a set of quality of service requirements that govern the resource allocation for the desired connection. Conventional allocation of network resources for connection management is by and large static. See column 1, lines 15-29. Static allocation of network resources is inefficient, and is inadequate for a dynamic networking environment where the user requirements and the quality characteristics of network resources are not static. See column 1, lines 4-45. Lee teaches that there is a need for a

system and method that addresses the problem of managing connections, and a connection-oriented integrated communications network, to maintain acceptable quality of service for the entire duration of each connection, to make quality of service improvements whenever it is possible, as well as to recover from connection failures and pre-emption. See column 2, lines 51-58.

Lee teaches the quality of service requirements for a connection specified in terms of performance constraints, resource constraints and priority constraints. A performance constraint may be non-negotiable or negotiable. A non-negotiable performance constraint is specified with an acceptable value that separates the entire range of possible performance values into two contiguous regions: unacceptable region, an acceptable region. A negotiable performance constraint is specified with a range of values bounded between an acceptable value and requested value. The requested value divide the acceptable region into two regions. The region bounded by the acceptable value together with the requested value is referred to as the agreeable region. The other region is referred to as the overkill region. During call set up, routing computes a path, and derives from an available value for each performance parameter. If the available value falls within the unacceptable region, the call is rejected. For non-negotiable performance constraints, if the available value falls within the acceptable region, the susceptible value is taken to be the agreed value. During a negotiation,

if the available value falls within the over kill region, a call is accepted. But only the requested value would be offered as the real value. See column 3, lines 43-column 4, line 20.

Resource constraints are specified with multi-level preferences but translated into acceptable and requested resource constraint sets. Each resource constraints set contains a subset of all resource attributes values. Resource constraints may also be directly specified in terms of these two constraint sets. During call set up, a routing first tries to define a path that satisfies the requested resource constraints set and other quality service constraints. If this attempt fails, the routing will use the acceptable resource constraint set for subsequent fall back path, and computations. When a path is accepted, the resource constraints set used in the final path computation is referred to as the effective resource constraint set. See column 4, lines 35-51.

Priority constraints are specified in terms of three connection priorities: establishment of priority, retention priority, and re-establish priority. Each priority may assume one of a predetermined number of levels during call. During call setup, pre-emption is permitted only after routing has tried in vain to find an acceptable path without resorting to pre-emption. See column 4, lines 52-58.

Lee teaches a novel approach to connection management by providing an enhancement procedure coupled with a free optimization procedure, update rules for negotiable performance constraints, and extension of the pre-emption mechanism. See column 4, lines 59-64. It is within this context, that the teachings of Lee must be considered.

Lee teaches when a call arrives it enters the establishment state. If there are not enough resources to support the call, the call will be rejected. Upon successful connection establishment, the call enters the information transfer state. When a call in the information transfer state is completed, it enters the release date, and the connection is subsequently taken down. When a call is in the information transfer state, its agreed performance values may fall out of match due to facility changes, link failures, and the like. As long as the quality of service degradation is within acceptable limits, the call remains in the information transfer state. If any degradation leads to a connection failure, or the call is pre-empted, the call enters the re-establishment state, so that the network may attempt to reestablish the connection. When a call is in the reestablishment state, the network attempts to find a new acceptable path for connection reestablishment. Upon successful reestablishment, the call reenters the information transfer state. The length of time in which a connection reestablishment attempt may be repeated is limited by a connection reestablishment delay. Beyond this delay, the reestablishment procedure is aborted, and the call enters the release date.

A call in the information transfer state is subject to enhancement to restore degraded performance values back to their agreed levels, and to avoid the use of resources that cannot satisfy the requested resource constraints. In the process of enhancement, the connection remains intact until a desired path is found. Enhancement involves finding a new path for the connection such that at least one agreed value is renegotiated to a more desirable value and none of the other agreed values is degraded. In addition, the requested resource constraints are also satisfied. See column 5, lines 5-58.

Lee teaches the enhancement procedure is coupled to a reoptimization procedure. Reoptimization is permitted only when the correctly agreed quality of service for each performance constraint is maintained at the most desirable level possible, subject to the appropriate update rules for the performance constraint. When a call is pre-empted, it enters the re-establishment state for reroute. There are three possible outcomes for connection reestablishment: hard pre-emption, soft pre-emption and disguise pre-emption. The reestablishment may fail, whereupon the call enters the release state and is subsequently disconnected. This outcome is referred to as a hard pre-emption. Should the pre-empted call be successfully reestablished, there are either not enough resources for the pre-empted call to be reestablished with its previously agreed throughput, or there are not enough resources such that its previously agreed throughput may be compromised. In the latter case, the reestablished call is said to have suffered soft pre-emption to make room for a higher priority

call. In the former case, since the connection has suffered but a minor glitch while it is being rerouted, this event is referred to as a disguised pre-emption. In both cases, the call returns to the information transfer state. Lee teaches for negotiable performance constraints, the susceptible and requested values may be updated for reestablishment, enhancement and reoptimization, according to one of the following rules: continuation rule, relaxation rule, and enhancement rule. In accordance with the continuation rule, both the requested and acceptable values associated with a given negotiable performance constraint are set to the previously agreed value, so that this value continues to be guaranteed. In accordance with the relaxation rules, the requested value is replaced by the previously agreed value. In accordance with the enhancement rule, the susceptible value is replaced by the previous agreed value. See column 6, lines 18-65.

Referring to Dighe, there is disclosed a UPC-based traffic control framework for ATM networks. Dighe teaches traffic classification is based on traffic characteristics of the offered load and quality of service requirements. The framework is based on the assumption that traffic management on a per-class basis is more realistic than traffic management on a per connection or per virtual channel basis. It is assumed that the service provider will define a set of m classes of service. Each class of service is targeted towards a set of specific applications, whose traffic characteristics and quality of service requirements from the network are similar. Associated with each class is a traffic parameter vector. The

traffic parameter represents the traffic characteristics of sources that subscribe to a given service class. Elements of the traffic parameter vector will be quantities such as the peak rate or burst size. Preferably, the traffic parameter vector is assumed to be the same as the traffic shaping vector defined by the ATM Forum. This is based on the dual-leaky bucket filter. See column 3, lines 12-43. It is within this context that all the teachings of Dighe are based on the dual-leaky bucket filter that the teachings of Dighe must be considered.

Dighe teaches a dual leaky bucket is define by three parameters namely: the peak rate, the sustained rate, and the compliance burst size. Once a source declares these parameters, the dual leaky bucket implementation of the UPC function at the UNI ensures that the source transmits no more than the compliant burst size of cells at the peak rate. If the burst size is larger than the compliant burst size, then the first compliant burst size cells are transmitted at the peak rate. The remaining cells in the burst are only transmitted at the sustained rate which can be at most equal to the peak rate. Hence, once the three parameters of the UPC function are known, the switch can determine a worst-case traffic profile for the offered traffic. Based on this worst-case traffic profile, the switch can decide to either accept the call, reject the call, or accept a call after suitably modifying the declared UPC parameters. See column 5, line 55-column 6, line 5. The remaining description found in Dighe focuses on the three parameters of the dual leaky bucket in regard to the specific architecture and equipment described by Dighe.

Applicant is pointing out this dual leaky bucket basis of the teachings of Dighe because of a criticalness of the system to be able to perform the dual leaky bucket implementation of the UPC function to be able to practice the teachings of Dighe.

The Examiner has concluded that the combination of Lee and Dighe arrive at applicants' claimed invention. Applicants respectfully submit that the teachings of Lee and Dighe cannot be combined because not only is there no teaching in the references themselves to combine them, but the teachings cannot be taken out of the context in which they are found, and by doing so, would be unworkable.

It is black letter patent law that there must be some teaching or suggestion in the references themselves to combine the references, and here, there is no such teaching or suggestion. Nowhere does Lee even consider the use of the UPC function, and this follows, because Lee focuses on a novel approach to connection management by providing an enhancement procedure coupled with a reoptimization procedure, update rules for negotiable performance constraints, an extension of the pre-emption mechanism. See column 4, lines 59-64. There is no need for UPC in Lee and it is not considered. In fact, there is no reason why one skilled in the art would even apply the UPC function (if they even could) to the teachings of Lee. Similarly, the dual leaky bucket that is the basis of the UPC function taught by Dighe has no need for the novel approach to connection management involving an enhancement

procedure coupled with a reoptimization procedure, update rules for negotiable performance constraints, and extension of the pre-emption mechanism, as taught by Lee. This also follows, since Dighe is not interested in any such enhancement procedure coupled with the reoptimization procedure and has no need for the same, since there is no place for the same in the operation of the system taught by Dighe and the dual leaky bucket. Applicants submit that the only reason that the Examiner is combining the teachings of Lee and Dighe is from the hindsight of applicants' claimed invention, and this is contrary to patent law. The Examiner cannot use applicants' claims as a road map to find the elements and limitations of applicants' claims in various references, and having found them, concluding that applicants' claimed invention is arrived at.

Furthermore, not only is there no teaching to combine Lee and Dighe, but when teachings from different references are combined, they must be done so in the context in which they are found. Applicants have just explained above the specific context of each of the references. The context of Lee is in regard to a novel approach to connection management providing an enhancement procedure coupled with the reoptimization procedure, and all the associated teachings and operation of Lee follows. The context of Dighe involves the UPC function that is based on the dual leaky bucket. These contexts have nothing to do with each other. The parameters that are required for the dual leaky bucket, and are identified by Dighe cannot be found and are not identified at all in the teachings of Lee. To combine the teachings

of the UPC that the Examiner suggests is found in Dighe, requires how the UPC is performed in Dighe must also be applied under patent law to the teachings of Lee. By doing so, not only does it not make any sense to apply the dual leaky bucket somehow or other to the operation of Lee, but it would require Lee to be completely redesigned and require significant amount of experimentation and testing to even possibly be workable. This only further supports the non obviousness of applicants' claimed invention. Accordingly for this reason, applicants' claimed invention is not obvious from, and is patentable over Lee in view of Dighe.

Moreover, precisely because of how different the contexts of each of these references are that one relies completely on the dual leaky bucket, and the other one has nothing at all to do with the dual leaky bucket but applies a totally different approach to connection management, one skilled in the art cannot combine these references and make the result operational.

Claims 1 and 8 include the limitation that the user specifies, for each quality of service requirement for which there is the entry in a UPC associated with the circuits, which requirements can be changed, and what order and by how much. There is no teaching or suggestion of this limitation in the applied art record.

The Examiner states that Lee discloses an apparatus for establishing circuits in an ATM network comprising a controller for which a user specifies every quality of service requirement of the circuit which can be changed, in what order, and by how much. The Examiner refers to column 3, lines 58-67, column 4, lines 1, 2 and 35-58 and column 6, line 64-67 and column 7, lines 1 and 2.

In regard to the limitation of the index found in Claim 1 of applicants, the Examiner simply states that the Examiner takes official notice that use of controllers and memories with indexes in ATM switches are conventional in the art. Therefore, it would be obvious to a person with ordinary skill in the art at the time of the invention to implement the invention of Lee using a controller and a memory with an index. Applicants respectfully traverse this conclusion.

As explained above, the Examiner is simply stating that since each of the elements of applicants' invention of Claim 1 already exist, then it would be obvious to take all these elements that are separate and disjoint, and conclude that applicants' invention of Claim 1 is arrived at. However, to reiterate, applicants are not suggesting that they discovered controllers or discovered memories or indexes, but they are claiming that they have invented a unique and novel way to use controllers and memories and indexes, as set out in the invention of Claim 1 to arrive at a result that heretofore has never been attained.

Furthermore, by the Examiner taking the simple approach of recognizing through official notice the existence of controllers and memories and indexes in ATM networks, it precludes the specific analysis directed by law of looking at the context of the specific teaching, as well as the requirement that there must be teachings in the applied art of record to combine the various teachings the Examiner is relying upon to arrive at applicants' claimed invention.

By the Examiner's analysis, it is submitted that most patents would never have been granted since the specific elements of any claimed invention already existed. It should be noted that probably 99 percent of all patents issued, are based on existing elements that are combined in a unique and novel way. Accordingly, it is respectfully submitted that the simple statement that there exists a controller and a memory and an index in ATM switches fails to arrive at the limitation in Claim 1 that the memory includes an index having the different quality of service requirements ordered in terms of priority for the controller to choose when the control relaxes the original quality of service requirements and attempts to establish the circuit. This is a very specific application of an index directed to the ordering of service requirements in terms of priority for the purpose of relaxing the original quality of service requirements in regard to attempts to establish a the circuit, as found in Claim 1 of applicants.

In regard to Claim 1, there is the limitation that the user specifies to the controller which order the requirements can be changed. Furthermore, there is the limitation that the memory includes an index having different quality of service requirements ordered in terms of priority for the controller to choose when the controller relaxes the original quality of service requirements and attempts to establish the circuit. It is respectfully submitted that these citations by the Examiner in Lee did not show this limitation.

Referring to the citations the Examiner relies on in Lee, the Examiner in the Office Action states that Lee teaches the user specifies which order the circuits are established and relies on figure 10 #1004, which shows selecting the "Next" performance parameter, which shows that the parameters are ordered. See page 2, paragraph 3 of the Office Action. Referring to figure 10 and specifically No. 1004, this is in regard to the steps of re-establishment as taught by Lee. As mentioned above, there are three connection priorities: the establishment priority, retention priority, and re-establishment priority. See column 4, lines 52-55. The re-establishment state only occurs if there is a degradation that leads to a connection failure, or the call is pre-empted. The call then enters the re-establishment state so that the network may attempt to re-establish the connection. See column 5, lines 15-24. Thus, Lee only teaches selecting a next predetermined performance parameter in the re-establishment state. Lee is silent about such a step in the establishment state and only teaches such a step in the reestablishment state. Thus, Lee is fully knowledgeable of this technique,

but failed to suggest or teach to apply such a step in the establishment stage. This indicates that Lee specifically does not teach such a step in the establishment state. The fact that Lee specifically teaches in the re-establishment state to select a next predetermined performance parameter but fails to suggest to teach such a step in regard to the establishment state, shows that the could have applied such a step since Lee knew about such a step, but specifically chose not to apply it to the establishment state.

In fact, in regard to the establishment state, Lee fails to teach or suggest the limitation of automatically relaxing the original quality of service requirements associated with a circuit for the circuit to be formed in the ATM network with available resources of the ATM network, as found in Claim 1 of applicants. In regard to the establishment state, Lee simply teaches that when a call arrives, it enters the establishment state. If there are not enough resources to support the call, there will be a failure to establish and the call will be rejected. Otherwise, upon successful connection establishment, the call enters the information transfer state. See column 5, lines 4-12.

It is submitted, the Examiner is taking the teachings of reestablishment and applying them to the establishment state in regard to Lee. However, Lee does not teach or suggest what the Examiner is suggesting, and the Examiner is in fact applying teachings from Lee to the establishment state, where they are not taught. Accordingly, applicants respectfully

disagree with the Examiner's statement that Lee teaches that the user specifies the order that the requirements can be changed to establish the circuit (to reiterate, to establish, not to reestablish the circuit).

With respect to column 3, lines 58-67, it simply states that the performance constraint may be non-negotiable or negotiable. A non-negotiable performance constraint is specified with an acceptable value that separates the entire range of possible performance values into two contiguous regions: unacceptable region, and acceptable region. A negotiable performance constraint is specified with a range of values bounded between an acceptable value and requested value. The requested value divides the acceptable region into two regions. The region bounded by the acceptable value together with the requested value is referred to as the agreeable region. The other region is referred to as the overkill region. As is clear, this language does not teach or suggest anything about having an index with different quality of service requirements ordered in terms of priority, let alone specifying the order in which the quality of service requirements are changed, as found in Claim 1.

In regard to lines 35-58 of column 4, it states resource constraints are specified with multi-level preferences, but are translated into acceptable and requested resource constraints sets. Each resource constraint set contains a subset of all resource attributes values, such as the requested resource constraints set is a subset of the acceptable resource

constraint set. Resource constraints may also be directly specified in terms of these two constraint sets. During call set up, routing first tries to find a path that satisfies the requested resource constraint and other quality service constraints. If this attempt fails, routing will use the acceptable resource constraint set for subsequent fall back path computations. When a path is accepted, the resource constraints set used in the final path computation referred to as the effective resource constraint set.

There is nothing in this language that sets out the order in which the quality of service requirements are chosen to be relaxed. This language specifically simply states that there are multi level preferences. In regard to any given constraint, they are specified into preferences, but this does not refer to any type of index, nor specifically the order in which the resource constraints are first chosen to then determine for that given resource constraint what multi-level preferences applicable. This is what is completely missing from the teachings of Lee in regard to the limitations of Claim 1.

It is submitted the Examiner is reading limitations of applicants' Claim 1 into the teachings of Lee, where these specific limitations are not present. In regard to lines 64-67 of column 6, and lines 1 and 2 of column 7, it simply states that in accordance with the enhancement rule of the present invention, the susceptible value is replaced by the previously

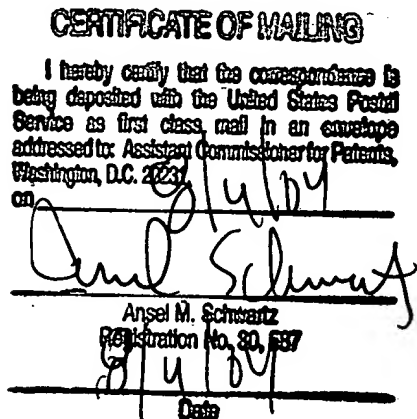
agreed value. It does not change the requested acceptable values that originally configured by the user. Accordingly, this citation adds nothing further to the limitations of Claim 1.

Accordingly, Claims 1, 2, 4-6, 8 and 9 are patentable over Lee in view of Dighe.

The Examiner has rejected Claims 7, 14 and 15 as being unpatentable over Lee in view of Dighe and further in view of Burns. Applicants respectfully traverse this rejection. Burns simply teaches the use of the SPVCs. However, Burns does not teach their use in regard to relaxing automatically equality of service requirements of a UPC associated with circuits. There is nothing to suggest that the same could occur and be used. Again, the Examiner is using hindsight to find the different elements of applicants' claims in different references, and having found the different elements in the different references concludes applicants' claims are arrived at. This is contrary to the law. Once again, the Examiner is simply saying that because there would be some advantage to combine the references, it is proper to combine them. The only reason it is advantageous to combine the references is from the teachings of applicants' own claims and specification, which is using hindsight. Accordingly, Claims 7 and 14, which are dependent to Claims 1 and 8, respectively, are patentable for the reasons that the parent claims are patentable, and in addition for their own limitations.

The Examiner has rejected Claims 12-14 as being unpatentable over Lee. For the reasons explained above, Lee does not teach or suggest anything at all in regard to the UPC function or an index. Accordingly, Claims 12-14, which are dependent to parent Claim 8, are patentable over Lee.

In view of the foregoing remarks, it is respectfully requested that the outstanding rejections and objections to this application be reconsidered and withdrawn, and Claims 1, 2, 4-9 and 12-17, now in this application be allowed.



Respectfully submitted,

MARK HELD, ET AL.

By Ansel M. Schwartz
Ansel M. Schwartz, Esquire
Reg. No. 30,587
One Sterling Plaza
201 N. Craig Street
Suite 304
Pittsburgh, PA 15213
(412) 621-9222

Attorney for Applicants